
Homogenization of Existing Inertial-Based Datasets to Support Human Activity Recognition

Hamza Amrani, Daniela Micucci, Marco Mobilio, Paolo Napoletano

Department of Informatics, Systems and Communication

University of Milano - Bicocca, Milan, ITALY

{hamza.amrani, daniela.micucci, marco.mobilio, paolo.napoletano}@unimib.it

Abstract

Several techniques have been proposed to address the problem of recognizing activities of daily living from signals. Deep learning techniques applied to inertial signals have proven to be effective, achieving significant classification accuracy. Recently, research in human activity recognition (HAR) models has been almost totally model-centric.

It has been proven that the number of training samples and their quality are critical for obtaining deep learning models that both perform well independently of their architecture, and that are more robust to intraclass variability and interclass similarity. Unfortunately, publicly available datasets do not always contain high quality data and a sufficiently large and diverse number of samples (e.g., number of subjects, type of activity performed, and duration of trials) Furthermore, datasets are heterogeneous among them and therefore cannot be trivially combined to obtain a larger set.

The final aim of our work is the definition and implementation of a platform that integrates datasets of inertial signals in order to make available to the scientific community large datasets of homogeneous signals, enriched, when possible, with context information (e.g., characteristics of the subjects and device position). The main focus of our platform is to emphasise data quality, which is essential for training efficient models.

1 Introduction

Human Activity Recognition (HAR) aims at automatically classifying activities performed by humans (including falls) by analyzing signals acquired by sensors [1, 2, 3]. Recent methods and approaches mostly exploit inertial sensors embedded in smartphones, smartwatches, fitness trackers, and ad-hoc wearable devices.

In recent years, deep learning techniques have been successfully applied for 1D signals, exploiting their capability to overcome most of the issues raised by traditional machine learning techniques, thanks to their properties of local dependency and scale invariance [4]. While deep learning methods are powerful and achieve high performance, they rely on very complex models that depend on estimating a large number of parameters, which in turn requires a considerable amount of available data [5], whose quality increases the performance of the classification process.

Building an effective dataset is a complex task. Several factors undermine its goodness. These include the naturalness with which users perform the tasks, the position of the device, the balance of subjects involved, the number of samples recorded, and so on. Even if the dataset design is done in a rigorous way, there are factors that unfortunately are not controllable: the intraclass variability and the interclass similarity. The former means that different people perform the same activity in different

ways, so a bijective association between signal and activity performed does not exist; the latter means that fundamentally different classes show very similar characteristics in terms of sensor data [3, 6, 7].

In the literature, a number of HAR datasets are available [8, 9, 10]. However these datasets, besides suffering from the above problems, are heterogeneous with each other. For example, signals are sometimes expressed in different units of measurement, may have different acquisition frequencies, and accelerations may include gravity or not. Moreover, ADLs and therefore the labels in the datasets, do not have a common dictionary or ontology, which leads in having similar labels that actually have different meanings. Thus, datasets cannot be used together without a significant effort to harmonise them.

The availability of a dataset containing a large number of samples, also obtained from the integration of existing datasets, is a well-known issue both in the field of ADLs recognition from inertial sensors and in other domains, such as that related to image processing [11]. In the context of ADLs recognition from inertial signals, Bartlett et al. proposed labels aggregation at the semantic level [12]. Labels of six existent datasets, resampled to 200 Hz using linear interpolation, were relabelled manually to reflect their semantic similarity, obtaining 13 different activity labels. The proposal however does not seem to consider details such as units of measurement or the presence of gravity in accelerations. Furthermore, the proposal allows only one configuration at 200 Hz. Recently, Siirtola et al. proposed a Matlab tool called Open HAR [13] that aggregates labels at a syntactic level but fails to consider the semantics of the original signals. Obinikpo et al. proposed a system for big data-d-health integration [14]. They split the integration into different layers: data acquisition, data processing, analytics, and application. Nevertheless, the proposal is too general concerning the homogenization of different data sources since its major concern is handling missing values while integrating databases.

The main contribution of this paper is the definition of a homogenization procedure that allows to integrate heterogeneous datasets in order to obtain a larger dataset to be used for the definition of recognition techniques. The procedure has been implemented and integrated in a platform termed Continuous Learning Platform (CLP) that makes available (i) a large amount of labelled inertial signals related to ADLs and falls; (ii) a catalogue of downloadable activity recognition models, and (iii) a service that, given a set of raw data, identifies the corresponding ADL. The platform is available at the following URL: <https://gitlab.com/Pervasive-Healthcare/CLP>.

The paper is organised as follows: Section 2 describes the homogenization procedure we identified; Section 3 provides an overview of the CLP platform and some implementation details; while Section 4 presents final remarks.

2 Homogenization Procedure

Homogenization involves two different phases: *signals homogenization* and *labels homogenization*.

Signals Homogenization. The signals homogenization procedure, as the name suggests, focuses on signals and handles three types of inconsistencies: differences in sampling frequencies, discrepancies in units of measurement, and the presence of gravity for acceleration signals. The process involves three steps:

- *Frequency uniformation.* A resampling operation takes place intending to modify the frequency of the time series. There are two different types of resampling: upsampling, when the sampling rate is increased compared to the original one, and downsampling, when the sampling rate is decreased compared to the original one. The number of samples obtained with the new frequency can be calculated starting from Equation 1.

$$numSamples = \frac{numOriginalSamples * newFreq}{originalFreq} \quad (1)$$

Once the new time series with the desired frequency has been obtained, the respective timestamp is calculated for each sample starting from 0. Information regarding the original sampling rate of the dataset is obtained by the metadata if available, or, if this information is not provided, it is calculated using the timestamp associated with the inertial data. When the sampling frequency has been obtained or estimated, the data can be resampled. For all type of sensors, we adopted a frequency of 50Hz. Literature suggests that about 50Hz is a suitable sampling rate that permits to model human activities [15].

- *Unit of measurement uniformation.* The unit of measurement, unlike the sampling frequency, must always be provided. It is not possible to trace this information from inertial data. The unit of measurement of a given sensor is converted to the desired one by a specific formula. For example, to convert an accelerometer from g to m/s^2 , it is necessary to add the gravitational acceleration, that is, multiplying by the gravitational acceleration constant, equal to $9.80665m/s^2$. For accelerometer, gyroscope, and orientation we adopted m/s^2 , rad/s , and $microTesla$ respectively.
- *Gravity uniformation.* To remove gravity and reduce noise or artefacts, a Butterworth filter is commonly applied [16, 17]. We considered a fifth-order 0.5 Hz low pass Butterworth filter with a Nyquist frequency of 25Hz, assuming the gravity force to have only low-frequency components. Since the information about gravity being included in the signal or not is often omitted, we decided to apply the Butterworth filter to all raw inertial signals.

Labels Homogenization. The label homogenization procedure aligns the labels of the ADLs in the dataset to be homogenized with the labels chosen as the reference ones. Because there is no shared definition for each ADLs and each dataset can have different and conflicting labels, relying on the syntactic similarity may not be sufficient (e.g., walk and walking describe the same activity although the two labels are syntactically different, on the other hand, sitting, and sitting down, may be syntactically similar, but they refer to different actions: being sit and actively sitting from a standing position). We provide two different approaches to make an accurate mapping:

- *Label Syntax Similarity (LSS).* This index indicates the syntactic similarity of two strings (labels) without considering the signal component. We used the Levenshtein distance, which is the minimum number of elementary modifications (deleting a character, replacing one character with another, or inserting a character) that allows transforming a string A into another string B. For example, the Levenshtein distance between walk and walking equals 3. This distance can be helpful in some cases, but it cannot provide adequate information to automatically choose the correct mapping.
- *Label Signal Similarity Distance (LSSD).* Considering the signal component from the time series, we extracted 21 features from the magnitude component of the signal for each window. These features describe the different properties of the signal both in the time and frequency domain. The average of the features extracted from the windows of a specific activity represents the entire ADL. Each dataset has a feature vector per ADL. The Euclidean distance is then applied to determine possible associations. If the distance of two activities is low, it suggests the associated signals are very similar, and the suggestion is a mapping between the two ADLs.
- *Magnitude comparison graph.* The procedure also includes a manual comparison of the magnitude of different labels' time series for each pair that satisfies the minimum equality criteria. A random time series is taken for each label to create the graph. This is a visual aid to help a user in deciding which mappings to perform.

The final decision about mappings is always left to the users. However, LSS and the LSSD are useful Decision Support Systems that ease the process and reduce the confusion that may arise from heterogenous labels.

3 Continuous Learning Platform

Continuous Learning Platform (CLP) is a freely available platform that implements the homogenization procedure described in Section 2 and provides the tools to both integrate new datasets and to retrieve homogenized datasets and recognition models.

Three main components constitutes CLP: *i) Data Collection*, that acquires a new dataset; *ii) Data Management* that homogenizes the new dataset and insert it in the incrementally built dataset; and *iii) Data Distribution*, which enables users and applications to query the platform and obtain homogeneous sets of labelled signals, but also ad-hoc trained classifiers.

Two more components complete the platform: the *Repository Manager* that deals with the management and internal storage of all the platform data, and a *Web Application* responsible for making all the services offered by the platform available through an intuitive graphical interface.

The individual components were developed in Python, while a message broker (RabbitMQ) has been used to permit efficient interaction between the components. The Web Application is built using the Angular framework, while REST APIs have been used for communication. The components are implemented as a Flask server to manage the REST calls coming from the Web Application. The platform uses MongoDB as a database.

The rest of the section will give a more detailed view of the three main components.

Data Collection includes the following modules. The *Dataset Loader* allows users to physically upload datasets and store them in a local repository. The *Driver Loader* handles the upload of datasets drivers, which are scripts provided by the users allowing the platform to correctly interpret the data in the dataset. The *Importer* executes the actual import by running the uploaded driver over the uploaded dataset, thus standardizing it according to a specific structure shared by all the datasets imported. It is noteworthy that, at this stage, the content of the datasets are still heterogeneous, while they share a common structure.

Data Management reifies the homogenization procedure described in Section 2. It includes the following modules. The *Data aligner* homogenizes each signal in a dataset at a specific frequency and unit of measurement, according to a common configuration. It also applies a Butterworth filter to remove the gravity from accelerations and to reduce noise and other artefacts. The *Feature extractor* computes the magnitude and a set of hand-crafted features [18] from inertial signals to determine signal similarity. Finally, the *Label comparator* unifies the dataset's labels to include a standard unified set, also considering the signal similarities. This module is semi-automatic: it provides suggestions on the assignment of labels, but ultimately it is up to the end-user to decide whether or not to accept the suggestions.

Data Distribution consists of the following modules. The *Query builder* handles users' queries for homogenized data. The *Classifier builder* trains classifiers according to user queries. Thanks to the homogenization procedure, the models generalize well in terms of inter and intra-subject variability regardless of the fact that are trained on data coming from different datasets. The *Classifier deployer* distributes trained classifiers according to users' requests. The *Online classifier* provides online services related to classification: given a set of inertial signals, it provides information regarding the subject's activity.

4 Conclusions

Human Activity Recognition is a challenging and active research field and has seen rapid growth in the past few years. The lack of large datasets reduces exploitation of deep learning techniques as they usually require large amounts of data that exceeds the size of individual datasets.

We propose a platform (CLP) that enables the integration and the distribution of data coming from heterogeneous sources. The main components of CLP have been fully implemented and are available at <https://gitlab.com/Pervasive-Healthcare/CLP>, while some aspects, such as the Web application is still under active development.

Eight datasets available in the literature have been already integrated and used to train a convolutional neural network and a recurrent neural network. These preliminary tests confirm the improvement of performances by combining existing datasets from different smartphones and various contexts and will be the focus of future work.

Future directions include polishing the Web application and an intensive test of the overall platform with real-world applications. Relying on our previous work on personalization [19], we are currently working on designing a software component that interfaces the CLP platform and an Android application to develop personalized models.

References

- [1] Jorge L Reyes-Ortiz, Luca Oneto, Albert Sama, Xavier Parra, and Davide Anguita. Transition-aware human activity recognition using smartphones. *Neurocomputing*, 171:754–767, 2016.
- [2] Daniela Micucci, Marco Mobilio, Paolo Napoletano, and Francesco Tisato. Falls as anomalies? an experimental evaluation using smartphone accelerometer data. *Journal of Ambient*

Intelligence and Humanized Computing, 8(1):87–99, 2017.

- [3] Carlos Medrano, Raul Igual, Inmaculada Plaza, and Manuel Castro. Detecting falls as novelties in acceleration patterns acquired with smartphones. *PloS one*, 9(4):e94811, 2014.
- [4] Ming Zeng, Le T Nguyen, Bo Yu, Ole J Mengshoel, Jiang Zhu, Pang Wu, and Joy Zhang. Convolutional neural networks for human activity recognition using mobile sensors. In *Proceedings of the international Conference on Mobile Computing, Applications and Services (MobiCASE)*, 2014.
- [5] Simone Bianco, Remi Cadene, Luigi Celona, and Paolo Napoletano. Benchmark analysis of representative deep neural network architectures. *IEEE Access*, 6:64270–64277, 2018.
- [6] Nicholas D Lane, Ye Xu, Hong Lu, Shaohan Hu, Tanzeem Choudhury, Andrew T Campbell, and Feng Zhao. Enabling large-scale human activity inference on smartphones using community similarity networks (csn). In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*, 2011.
- [7] Christian Krupitzer, Timo Sztyler, Janick Edinger, Martin Breitbach, Heiner Stuckenschmidt, and Christian Becker. Hips do lie! a position-aware mobile fall detection system. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2018.
- [8] Eduardo Casilari, José-Antonio Santoyo-Ramón, and José-Manuel Cano-García. Analysis of public datasets for wearable fall detection systems. *Sensors*, 17(7):1513, 2017.
- [9] I. Cleland, M. P. Donnelly, C. D. Nugent, J. Hallberg, M. Espinilla, and M. Garcia-Constantino. Collection of a diverse, realistic and annotated dataset for wearable activity recognition. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, 2018.
- [10] Anindya Das Antar, Masud Ahmed, and Md Atiqur Rahman Ahad. Challenges in sensor-based human activity recognition and a comparative analysis of benchmark datasets: A review. In *Proceedings of the Joint International Conference on Informatics, Electronics Vision (ICIEV) and International Conference on Imaging, Vision Pattern Recognition (icIVPR)*, 2019.
- [11] Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. Snorkel: Rapid training data creation with weak supervision. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, 2017.
- [12] James Bartlett, Vinay Prabhu, and John Whaley. Actionnet: A dataset of human activity recognition using on-phone motion sensors. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.
- [13] Pekka Siirtola, Heli Koskimäki, and Juha Rönning. Openhar: A matlab toolbox for easy access to publicly open human activity data sets. In *Proceedings of the ACM International Joint Conference and International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers (UbiComp)*, 2018.
- [14] Alex Adim Obinikpo and Burak Kantarci. Big data aggregation in the case of heterogeneity: a feasibility study for digital health. *International Journal of Machine Learning and Cybernetics*, 10(10):2643–2655, 2019.
- [15] Nishkam Ravi, Nikhil Dandekar, Preetham Mysore, and Michael L. Littman. Activity recognition from accelerometer data. In *Proceedings of the 17th Conference on Innovative Applications of Artificial Intelligence*, 2005.
- [16] D Gordon E Robertson, Graham E Caldwell, Joseph Hamill, Gary Kamen, and Saunders Whittlesey. *Research methods in biomechanics*. Human kinetics, 2013.
- [17] Ryo Takeda, Giulia Lisco, Tadashi Fujisawa, Laura Gastaldi, Harukazu Tohyama, and Shigeru Tadano. Drift removal for improving the accuracy of gait parameters using wearable sensor systems. *Sensors*, 14(12):23230–23247, 2014.

- [18] Anna Ferrari, Daniela Micucci, Marco Mobilio, and Paolo Napoletano. Hand-crafted features vs residual networks for human activities recognition using accelerometer. In *Proceedings of the IEEE International Symposium on Consumer Technologies (ISCT)*, 2019.
- [19] Hamza Amrani, Daniela Micucci, and Paolo Napoletano. Personalized models in human activity recognition using deep learning. In *Proceedings of the International Conference on Pattern Recognition (ICPR)*, 2021.